

# StorNext 6.0.6 Release Notes

## Contents

What's New in StorNext 6.0.6 .....	2
StorNext Compatibility .....	6
Supported StorNext Upgrade Paths and Upgrade Considerations .....	7
Compatibility Between StorNext and Other Products .....	8
General Considerations .....	9
Upgrading Appliances .....	10
Appliance Release Notes .....	10
Known Issues .....	10
Contacting Quantum .....	22

---

© 2018 Quantum Corporation. All rights reserved. Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law. Artico, Be Certain (and the Q brackets design), DLT, DXi, DXi Accent, DXi V1000, DXi V2000, DXi V4000, DXiV-Series, FlexSpace, FlexSync, FlexTier, Lattus, the Q logo, the Q Quantum logo, Q-Cloud, Quantum (and the Q brackets design), the Quantum logo, Quantum Be Certain (and the Q brackets design), Quantum Vision, Scalar, StorageCare, StorNext, SuperLoader, Symform, the Symform logo (and design), vmPRO, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners. Products mentioned herein are for identification purposes only and may be registered trademarks or trademarks of their respective companies. All other brand names or trademarks are the property of their respective owners. Quantum specifications are subject to change.

---

# What's New in StorNext 6.0.6

## Purpose of this Release

The StorNext 6.0.6 release combines an array of appliance and StorNext specific software fixes with important functional enhancements. Packaged to be easily installed on StorNext appliances, this release adds support for LTO-8 with Linear Tape File System (LTFS) formatted media along with end-to-end support for the CentOS 7 update 4 distribution.

See [StorNext 6 Compatibility Guide](#) for complete details about supported configurations. Additionally, see the [Linear Tape File System \(LTFS\) Media Format](#) section for updated information about the use of LTO-8 tape devices with LTFS media.

## New Features and Enhancements in StorNext 6.0.6

This release includes the following new features and enhancements:

- Support for LTO-8 tape devices when used with Linear Tape File System (LTFS).
- Full support for the CentOS 7 update 4 distribution.

## Compatibility and Support

The [StorNext 6 Compatibility Guide](#) provides the basic compatibility for StorNext 6.0.6, includes the StorNext components supported, operating systems and service packs, libraries and drives, browsers, virtual machines, and appliance support. Listed below are just a few of the types of information available to you in the [StorNext 6 Compatibility Guide](#).

- **Upgrade Paths:** Provides information on what upgrades to this release are supported.
- **Appliance Support:** Provides information on what StorNext and Lattus appliances are supported with this release or are compatible with it.
- **Operating Systems and Platforms:** Provides information on what StorNext components run on various operating systems and service packs. Also includes which operating systems have been newly added or removed.
- **Client Interoperability:** Provides information on what StorNext clients running other versions of StorNext are compatible with metadata-controllers (MDCs) running this release.
- **Virtual Machine Support:** Provides information on what StorNext components running on selected operating systems and service packs are supported in virtual machines.
- **Compatibility with Other Products:** Provides information on references to additional StorNext sold-separately products that are supported with this release.
- **Browser Support:** Provides information on what versions of browsers are supported with the GUI in this release.

- **Drives and Libraries:** Provides information on what Quantum and 3rd party drives and libraries are supported with this release.

## Fixed Issues and Enhancements Addressed in StorNext 6.0.6

Operating System	Change Request Number	Service Request Number	Description
All	41413	1504258, 1635952, 3471000, 3526376, 3662584, 340173	mover should handle no sense check conditions better
All	44209	1570118, 1552386, 354261, 388885	Request to prevent the MDCs from SMITHing on startup if TSM fails to come up
All	48045	n/a	SELECT COUNT(*) FROM TIER000FILES1 makes showc -t painful
All	65522	357880	When nas_vip is configured for NAS cluster, the nas_vip should also update /usr/cvfs/config/dpserver file
All	69412	356235	VopNodeTokenChange threads waiting for a long running VopPunchHole trigger an OpHangLimit panic
All	69689	401102	hang in cvfs while submitting 0 length buffer
All	69741	400801	qbm-ladder requires perl DBD::SQLite.pm module
All	69772	404185	GUI: Stripe group actions: add stripe group only uses one LUN when multiple LUNs are selected
All	70062	373272	Inability to use mv on a file. Error produced from TSM.
All	70091	374701, 364890, 385126	Thread hang in open_check_busy() cause FSM OpHangLimitSecs (give-up for old clients)
All	70095	373633	snacl assumewinbind separator = +
All	70124	376032	MSM memory leak in LibManager
All	70148	372657	Object storage path remains IN USE after aborted fsfilecopy
All	70158	0373915, 404469	WSAR_Agent core dumps with segmentation fault
All	70159	378438, 374601	TSM not storing files from clients in some cases when time stamps are preserved.

Operating System	Change Request Number	Service Request Number	Description
All	70172	n/a	fsmpm debug log, nssdbg.out is filled with messages about vote call inhibited
All	70250	396019, 386306, 397004, 398663, 401382, 402587, 401544, 407324	Partial mdarchive backups sometimes fail because of a missing QRDB tablet file
All	70266	381540	SN605 Windows client sometimes do not mount fsforeign file systems if the fsnameservers file is empty
All	70298	381492	mdarchive_scan can cause FSM OpHangLimtSecs PANIC after upgrade to SN6
All	70307	387032, 385789	CCTL: get_cctls_from_host(): Received control info from hostspam in /var/log/messages
All	70325	383976, 399339, 397419, 398829	fs_policyd may core if mdarchive during SN6 managed FS conversion
All	70370	387837, 388093	File stores are failing because the file attributes change
All	70377	388885	Delayed FSM activation of managed file system caused TSM startup failure and then SMITH
All	70381	385126	hang since client and fsm repeatedly exchanging revoke and return messages
All	70416	388033	RYO SN6 Upgrade: install.stornext incorrectly select wrong path for mdarchive
All	70423	393905	Image activation of Java can trigger deadlock between cv_rwlock and mm_sem semaphore
All	70442	385126	cvadmin latencytest behaviour is highly confusing for active clients
All	70443	n/a	fs_altstore keeps old staging files in .AltStoreStagedir
All	70523	388885	FSM with metadata archive enabled unnecessarily calls stat(2) on every qdb file

Operating System	Change Request Number	Service Request Number	Description
All	70533	394300	sgmanage utilities should require central control administrative privilege for changing configuration
All	70580	397688	msg: Info: A .auth_secret file has been found filling message logs
All	70601	398046 398079	FSPM cannot join cluster if we use an .auth_secret file
All	70606	397969, 399081, 403292	FSM panics during SN 6 upgrade while building the mdarchive with active clients
All	70633	398922	SNFS 6.0.5 OpHangLimitSecs VOP-Class-0 Type-9 Subtype-13 (deadlock with open_lock, pclient lock, and sb_ClientWOpens)
All	70655	399230	Excessive gateway license RAS events used up all inodes on root filesystem
All	70667	398552	Cannot determine fileName or inode info for a given file
All	70682	399399	backup fails if metadata archive directory is not the default "/usr/adic/database/mdarchives/"
All	70786	399295	MSM memory leak in LibManager for vspoolqry and vsmedclassqry
All	70798	404467	Need an option to cvadmin to bypass host resolution.
All	70835	401251	nss_ctl_template generates an nss_ctl.xml file that does not allow any access to the file system
All	70839	399858	GUI regression. Tape drives that are offline not shown on dashboard page.
All	70840	401446	GUI "Report->Files" fails to display file info because of a java exception on encryption type
All	70990	406758	sgoffload can fail when the fsm REST interface becomes too slow and produces misleading error messages

Operating System	Change Request Number	Service Request Number	Description
All	71037	407630, 409786, 411650, 419565, 420464	Kernel panic in ExtApiXXXByName code path
Windows	68837	337136, 350237	Windows Exception in cvAcqWriteLock()

## StorNext Compatibility

For information on StorNext compatibility with operating systems, kernel versions, hardware platforms, drives, libraries, StorNext Appliances, StorNext client interoperability, and other compatibility items, see the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

**Note:** SNAPI and Partial File Retrieval information is provided in separate documents.

### Quantum Operating System Upgrade Support Policy

StorNext supports any security or functional bug update that applies to the current StorNext-supported Red Hat update level or SUSE Linux Service Patch. StorNext does not support upgrading to an update level or service patch beyond the currently supported levels shown in the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

### StorNext and Linux Interoperability

Newer versions of the Linux `tail` command leverage the `inotify` mechanisms within Linux. The `inotify` mechanisms in Linux are not triggered by file updates coming from other StorNext nodes.

When using the `tail` command on files located in StorNext, Quantum recommends using the following option:

```
---disable-inotify
```

Recommended usage:

```
tail ---disable-inotify -f filename
```

# Supported StorNext Upgrade Paths and Upgrade Considerations

## StorNext Software Upgrade Matrix

For information on which StorNext versions allow you to upgrade directly to this release, refer to the **StorNext Software Upgrade Matrix** section in the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

## Considerations for the StorNext File System Directories

On upgrades to StorNext 6.0.6, note that the attributes of many directories in the StorNext file system show much smaller sizes, even zero sizes, where these same directories showed non-zero sizes in previous releases of StorNext. This is expected behavior.


## Journal Size Guidelines

The absolute minimum Journal Size in StorNext 6.0.6 is 4 MB. If a file system is configured with a Journal Size smaller than 4 MB, the Journal Size must be increased prior to upgrading. The recommended Journal Size is 64 MB. New file systems must have a Journal Size of 64 MB or larger.

## Distributed Data Mover (DDM) Guidelines

Distributed Data Movers (DDMs) must be upgraded to the same version of StorNext that the Metadata Controller (MDC) is running.

---

 **WARNING:** Upgrades (such as platform, service pack, etc.) are intended to be done to all systems present in a given deployment. For example, if Xcellis, M660, M440, Pro Foundation, Artico, and G300 are present, they all must be upgraded. One appliance cannot be "left behind".

## Considerations When Upgrading NFS Server Nodes to StorNext 6.0.6

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 6.0.6, special consideration must be made for Linux NFS servers.

To prevent issues with mounted NFS clients, NFS clients must be unmounted prior to upgrading StorNext on the NFS server. If unmounting all NFS clients is not an option during the upgrade, Quantum suggests using the "compat32" mount option on NFS servers.

# Compatibility Between StorNext and Other Products

The following sections provide information regarding compatibility between this release and StorNext components and features.

## Appliance Controller

To view supported Appliance Controller software configurations, see the *Appliance Controller Compatibility Guide* available online at [http://www.quantum.com/acc\\_cg](http://www.quantum.com/acc_cg).

## Infiniband

StorNext 6.0.6 works with Infiniband SRP (SCSI RDMA Protocol) attached storage for Linux and Windows 2008R2.

## Lattus

See the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#) for information about compatibility between Lattus and StorNext 6.0.6.

---

**i Note:** Object Storage documentation is available online at <http://www.quantum.com/lattusdocs>.

## Partial File Retrieval

StorNext Partial File Retrieval (PFR) is a product which enables you to quickly retrieve and utilize segments of large media files, rather than the entire file, based on time-code parameters.

---

**i Note:** For Quantum Cloud Storage, PFR is not supported for copies with client-side encryption or compression. It is only supported for copies with server-side encryption or without encryption and compression.

For information about compatibility between PFR and StorNext 6.0.6, see the *StorNext Partial File Retrieval Compatibility Guide* in the [StorNext 6 Documentation Center](#).

## StorNext Web Services

StorNext Web Services enables you to run third-party application program interfaces (APIs) with StorNext. To view the latest commands supported by the StorNext Web Services, refer to the [StorNext 6 Web Services Guide](#) in the [StorNext 6 Documentation Center](#).



## Apple Xsan

Xsan is software that enables multiple Mac computers to concurrently access hundreds of terabytes of content on Xserve RAID or Promise RAID storage over high-speed Fibre Channel which allows you to share data faster and consolidate projects. Quantum supplements this solution with StorNext data management software, enabling Apple Xsan customers to use applications running on Windows, Linux, and UNIX with their Xsan and share content across more systems.

For information about compatibility between Apple Xsan and StorNext 6.0.6, refer to the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

## Supported Browsers

For information on browsers supported with the StorNext GUI for this release, refer to the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

For all other components and features, see the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

---

# General Considerations

This section provides information about items to consider for StorNext 6.0.6.

## Checksum Performance Considerations

**Note:** Generating MD5 checksums is a CPU-intensive operation.

Current StorNext metadata controller and Mover hardware is able to calculate MD5 checksums at around 300 MB/s to 500 MB/s. For newer generation tape technology, the maximum throughput might exceed the rate at which the system can generate checksums. In this case, the MD5 checksum calculation will define the throughput of a single data movement operation. With multiple movement streams, MD5 calculations will be done in parallel across the streams and aggregation of performance will be seen.

---

# Upgrading Appliances

For instructions on upgrading your firmware, see [Upgrade the System \(Upgrade Firmware\)](#) on the *Appliance InfoHub Documentation Center* ([www.quantum.com/ApplianceInfoHub](http://www.quantum.com/ApplianceInfoHub)).

---

## Appliance Release Notes

Refer to the respective Release Notes document for important information you should know about your system.

- [Xcellis Foundation](#)
- [aiWARE for Xcellis](#)
- [Xcellis Workflow Extender](#)
- [Xcellis Workflow Director](#)
- [Artico](#)
- [Pro Foundation](#)
- [G300](#)
- [M660](#)
- [M440](#)

---

## Known Issues

The following sections list known issues in this release of StorNext, as well as associated workarounds, where applicable:

- 
- **Note:** If you encounter one or more of the issues listed in this section, please contact Quantum Customer Support and report the issue(s) you encountered. Also inform the support representative whether you were able to successfully work around the issue(s) by using the provided workaround. Doing these things will help Quantum prioritize the order in which known issues are addressed in future StorNext releases.

### StorNext File System Known Issues

The table below lists known issues specific to the StorNext File System.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	54834	3505208, 3516356	<p>If a file is being copied to the StorNext file system using Windows Explorer and Windows Explorer crashes before it finishes copying all the data, the file might contain data blocks from old, deleted files. This problem occurs because Windows Explorer sets EOF to the size of the file before it writes the data to the file. This leaves a gap of uninitialized data in the file.</p> <p><b>i Note:</b> This problem can also occur with other programs that set EOF beyond the end of data.</p> <p>This problem does not occur if Windows Explorer encounters an error while writing the file; Windows Explorer will delete the partially written file.</p> <p><b>Workaround:</b></p> <p>To prevent this problem from occurring on StorNext, you can use the StorNext "client configuration" application's advanced mount option "Restrict Pre-allocation API" on Window systems and the "protect_alloc=yes" mount option on Linux systems. This option will set the unwritten parts of the file to zero. When this option is set, non-root users are unable to use the preallocation ioctl. This option also implies sparse=yes.</p> <p>For more information on this option, see the man page <code>mount_cvfs(8)</code>. The sparse option will introduce some overhead when using Windows Explorer. Before setting the <code>protect_alloc</code> option, see the sparse option in <code>mount_cvfs(8)</code> for a description of how it changes StorNext behavior.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
------------------	-----------------------	------------------------	------------------------

All	70306	n/a	The <b>snrecover</b> utility only recovers files and directories that have been previously deleted. That implies that it cannot be used to find and recover some or all deleted contents within an existing directory. For example, the following removes all contents from directory <b>bar</b> and then attempts to recover the deleted contents:
-----	-------	-----	---

```
# rm -rf /stornext/snfs1/bar/*
# ls -l /stornext/snfs1
total 0
drwxr-xr-x 4 root root 0 Nov 13 13:34 bar
drwxr-xr-x 2 root root 0 Nov 13 13:34 lost+found
# snrecover -p /stornext/snfs1/bar -t 2017-12-03:04:05:06
```

In the example, **snrecover** is attempting to recover a directory that exists (**bar**). This is not a valid **snrecover** request so **snrecover** will return an error. If a directory exists, then it cannot be recovered. The deleted content of directory **bar** can be recovered using successive calls to **snrecover**, one for each deleted file or directory.

**Workaround:**

There is no workaround for this issue, as the **snrecover** utility is working as designed.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	71147	n/a	<p>If a client is not directly connected to the metadata network, then the file systems are not visible to that client. As a result, the client does not have access to mount any file systems.</p> <p>A client is directly connected to the metadata network when one of its interfaces is in the same subnet as the metadata network. The metadata network is defined by the addresses of the name servers and is configured within the <b>fsnameservers</b> file.</p> <p><b>Workaround:</b></p> <p>To correct this, use the <b>foreign servers</b> feature of StorNext to locate file system services. This feature allows a client to connect directly to a file system service by placing the address of the MDC (not the name service, although they can be the same) in the <b>fsforeignservers</b> file.</p> <p>See <a href="#">Configuration &gt; Name Servers</a>, or the <b>fsforeignservers</b> command in the <a href="#">StorNext 6 Man Pages Reference Guide</a> for more information.</p> <p><b>i Note:</b> The <b>foreign servers</b> feature is supported on all StorNext client platforms.</p>
Windows	69366	n/a	<p>After installing the StorNext client package on a Windows system, StorNext does not start. If you try to start StorNext, the following error message appears:</p> <pre>Error 2: 'Error starting Windows Service 'cvfsfilter'</pre> <p>You might encounter this problem when installing a StorNext 6.0 client-only package under the following circumstances:</p> <ul style="list-style-type: none"> <li>• When you remove the StorNext file system before installing the StorNext client package.</li> <li>• When you upgrade a StorNext client package.</li> </ul> <p><b>Workaround:</b></p> <p>To correct this, reboot the Windows system and reinstall StorNext.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
Windows	70190	n/a	<p>If you are using a Windows MDC, you have one or more StorNext file systems configured in HA mode (<b>haFsType</b> set to <b>HaUnmanaged</b>) and do not have a valid <b>ha_peer</b> file, then any file systems configure in HA mode will fail to start when the MDC is upgraded to StorNext 6.0.5. StorNext 6.0.5 enforces having a valid <b>ha_peer</b> file for HA configured file systems.</p> <p><b>Workaround:</b></p> <p>To correct this, you can either create a valid <b>ha_peer</b> file or disable HA mode. For more information on configuring HA on a Windows MDC, see <a href="#">Configuring a High Availability System on Windows</a>.</p>
Mac OS	67871	n/a	<p>macOS releases 10.12 through 10.12.6 contain an issue in which the operating system crashes and reboots if a rename is performed on a file inside a managed directory.</p> <p><b>Note:</b> The issue only occurs if rename tracking is <b>disabled</b> on the file system.</p> <p><b>Workaround</b></p> <p>To prevent this issue, <b>enable</b> rename tracking on managed file systems with Xsan clients.</p>
Mac OS	66948	322824, 336945	<p>If you access StorNext file systems from Apple Xsan clients, then you might encounter I/O error messages in the system log that do not contain details about real I/O errors detected on the Xsan client.</p> <p><b>Workaround</b></p> <p>If you encounter the errors on an Xsan client, contact Apple.</p>

## StorNext Storage Manager Known Issues

The table below lists known issues specific to StorNext Storage Manager.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	43320	1581004	<p>File retrieves from media to disk can be suboptimal for fast tape drives like the Oracle STK T10K drives. This scenario can occur when the retrieve event is initiated on a host that is different from the host running the mover process, which requires the use of synchronous direct I/O.</p> <p><b>Workaround:</b></p> <p>To work around this issue and achieve optimal performance for both file stores and retrieves with the T10K drives, increase the default I/O size used by the mover process and make the mover process use asynchronous buffered I/O when the use of synchronous direct I/O is not required, using the following steps:</p> <ul style="list-style-type: none"> <li><b>i Note:</b> This workaround might also help improve the performance of the faster LTO drives like LTO-6, and LTO-7 by updating the <b>FS_LTO_BLOCK_FACTORsysparm</b>.</li> <li><b>i Note:</b> Changes to <b>FS_xxx_BLOCK_FACTOR</b> only affects tapes formatted after the change.</li> </ul> <ol style="list-style-type: none"> <li>1. Change the <b>FS_T10K_BLOCK_FACTORsysparm</b> from 8 to 32 by adding the following entry to <code>/usr/adic/TSM/config/fs_sysparm_override</code>: <pre style="background-color: #f0f0f0; padding: 5px;">FS_T10K_BLOCK_FACTOR=32;</pre> <ul style="list-style-type: none"> <li><b>i Note:</b> The T10K default I/O block size is 512 KB or 8 * 64 KB. With the block factor changed to 32, the new T10K I/O block size will be 2 MB or 32 * 64 KB. Presently, the <b>FS_T10K_BLOCK_FACTORsysparm</b> must not be set to a value that exceeds 32.</li> </ul> </li> <li>2. Restart Storage Manager to ensure the change in <b>Step 1</b> goes into effect: <pre style="background-color: #f0f0f0; padding: 5px;"># tsmstop # tsmstart</pre> </li> <li>3. Verify the <b>FS_T10K_BLOCK_FACTORsysparm</b> contains the new value:</li> </ol>

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<pre data-bbox="683 331 1458 436"># showsysparm FS_T10K_BLOCK_FACTOR FS_T10K_BLOCK_FACTOR=32</pre> <ol style="list-style-type: none"> <li data-bbox="618 478 1458 548">4. Save the current copies of your <code>/etc/fstab</code> on the MDCs and the DDM clients.</li> <li data-bbox="618 569 1458 674">5. Modify <code>/etc/fstab</code> on the MDCs and the DDM clients to use the <code>auto_dma_write_length</code> and <code>auto_dma_read_length</code> mount options as follows: <pre data-bbox="683 709 1458 856">snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_dma_read_length=16m 0 0</pre> </li> <li data-bbox="618 898 1149 926">6. Unmount and re-mount your file systems.</li> <li data-bbox="618 947 1458 1192">7. Use new T10K media to store a copy of the file from the disk. <div data-bbox="683 989 1458 1192"> <p><b>i Note:</b> <b>Step 7</b> is very important; when the new copy is made to the new tapes, the new tapes are labeled with a 2 MB block size, which is used for subsequent writes or reads to and from the media. Tapes on which <b>fsformat</b> was run before the change will use the block factor in use at that time. This change will not impact those tapes.</p> </div> </li> </ol>
All	46693	n/a	<p>Executing the command <b>snbackup -s</b> while a full or partial backup is running might result in a message that <code>/usr/adic/TSM/internal/locks/backup.1f</code> is in an invalid format.</p> <p>This is due to the <b>snbackup -s</b> process reading the <b>backup.1f</b> status file while the backup process is updating it.</p> <p><b>Workaround:</b></p> <p>Ignore the message; to clear-up the process, re-execute the command <b>snbackup -s</b> (provided that the backup is not writing to the <b>backup.1f</b> status file while <b>snbackup -s</b> is trying to read it again).</p>



Operating System	Change Request Number	Service Request Number	Description/Workaround
All	47833	n/a	<p>When copying files between media using the CLI command <b>fsmedcopy</b>, the file is not re-segmented to match the segment size of the destination media. Rather, the original segments are copied to the target media type and the distribution of segments across destination media will, therefore, be the same as the distribution on the source media.</p> <p><b>i Note:</b> This behavior might cause file data segment distribution to be sub-optimal on the destination media.</p> <p><b>Workaround:</b></p> <p>Currently, a workaround does not exist for this known issue.</p>
All	69265	n/a	<p>Your DDMs might experience a timeout if you try to connect to the database. The issue is identified by an error log in <b>/usr/adic/TSM/logs/tac</b> which contains the text:</p> <pre>Process fs_moverd on &lt;host&gt; timed out trying to connect to the database. This usually indicates network connectivity trouble. Try increasing the timeout value by setting the connect_timeout value in /usr/adic/mysql/my.cnf. The default setting is 10 seconds so the new value should be larger.</pre> <p><b>Workaround:</b></p> <ol style="list-style-type: none"> <li>1. Increase the database connection timeout value by adding the following line to <b>/usr/adic/mysql/my.cnf</b> under the section labeled <b>[mysqld] connect-timeout=240</b>.</li> <li>2. Cycle the Storage Manager in order to pick up the updated timeout value.</li> </ol>
All	69341	n/a	<p>If you have the IBM APFO driver installed and configured, then when you perform an <b>fsmedread</b> operation of a partial tape block from a full tape block, the operation can fail with <b>errno=12</b>.</p> <p><b>i Note:</b> This issue affects all IBM APFO versions 3.0.19 and earlier, and has an impact primarily on disaster recovery procedures.</p> <p><b>Workaround:</b></p> <p>To correct this, perform an <b>fsmedread</b> operation without the IBM APFO driver.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	70991	n/a	<p>After importing an LTFS tape whose data was not originally written by StorNext, a multi-file retrieve command successfully retrieves the first file in the list, but fails to retrieve the remaining files.</p> <ul style="list-style-type: none"> <li>This issue only affects LTFS tapes whose data was NOT originally written by StorNext. Tapes that were written by StorNext and subsequently exported and re-imported are not affected by this issue.</li> <li>This issue only affects multi-file retrieves (including recursive retrieves). Single-file retrieves (including event-driven retrieves) are not affected.</li> <li>This issue only affects files that were imported with the "media ingest" function of <b>fsimport</b>. Any files imported with the "file ingest" function are not affected because they are imported in whole to the disk file system, and the media is not kept in the system.</li> </ul> <p><b>Workaround:</b></p> <ul style="list-style-type: none"> <li>For files affected by this issue, retrieve files one at a time.</li> <li>When importing non-StorNext LTFS tape data, consider using the "file ingest" function of <b>fsimport</b>. This function ingests the tape data, but not the tape itself. The tape is ejected from the system upon completion of the <b>fsimport</b> operation.</li> </ul>

## StorNext GUI Known Issues

The table below lists known issues specific to the StorNext GUI.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	69360	n/a	<p>Using <b>autofs</b> to mount a StorNext file system on an MDC is not supported when the same file system also has a native mount point.</p> <p>For example, if the StorNext file system <b>snfs1</b> is mounted as <b>/stornext/snfs1</b>, then the MDC should not also have an <b>autofs</b> configuration that mounts it on the MDC in another location such as <b>/space/snfs1</b>. Doing so, causes the <b>fsCheckAffinities</b> and <b>fsCheckTsmFilesystemConfig</b> health checks to fail and generate RAS tickets.</p> <p>Additionally, this might cause the StorNext GUI to fail unexpectedly for certain operations.</p> <p><b>Workaround:</b></p> <p>There is currently no workaround for this issue. If you experience this issue, contact <a href="#">Quantum Technical Support</a>.</p>

---

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	69958	373823	<p>There is a known issue where stripe group expansion using the StorNext GUI can fail and requires manual intervention to restore file system operation.</p> <p>Stripe group expansion allows an additional disk (LUN) to be added to an existing stripe group, growing the file system user data space.</p> <p>The other way to add space to a file system is to create a new stripe group and add this to the file system.</p> <p>When the GUI executes stripe group expansion, it stops the file system, modifies the configuration file and runs <b>cvupdatefs</b>. If <b>cvupdatefs</b> takes more than 5 minutes to complete, the GUI kills the <b>cvupdatefs</b> utility and reports an error. At this point the file system does not start because the configuration file does not match the current state of the metadata.</p> <p><b>Workaround:</b></p> <p>To address this issue, you can do one of two things:</p> <ul style="list-style-type: none"> <li>• Optimally, you would not attempt the stripe group expansion at all. Instead add a new stripe group to the file system.</li> <li>• If stripe group expansion is deemed necessary, use the <b>cvupdatefs</b> CLI directly instead of through the GUI. This method does not have any time limitations.</li> </ul> <p>In the case that the StorNext GUI attempt was made and hit the 5 minute timeout, file system operations can be resumed by performing the following:</p> <ol style="list-style-type: none"> <li>1. Restore the previous version of the configuration file. This can be found in the following directory: <div data-bbox="695 1360 1242 1394" data-label="Text" style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <pre>/usr/cvfs/data/&lt;fs&gt;/config_history</pre> </div> </li> <li>2. Run the <b>cvfsck</b> utility to verify and potentially correct the metadata.</li> <li>3. Start the file system.</li> </ol>

Operating System	Change Request Number	Service Request Number	Description/Workaround
Linux	47954	n/a	<p>The Safari browser becomes unresponsive when you attempt to configure an Email server using the StorNext GUI.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, perform the following procedure:</p> <ol style="list-style-type: none"> <li>1. Shut down the Safari browser window(s).</li> <li>2. Restart the Safari browser, and then retry the operation.</li> <li>3. Uncheck the <b>Verify SMTP Server Connectivity</b> box, and then retry the operation.</li> <li>4. Set <b>Authentication</b> to <b>NONE</b>, and then retry the operation.</li> <li>5. Disable the Safari <b>User names and passwords</b> AutoFill under <b>Safari &gt; Preferences &gt; AutoFill</b>, and then retry operation.</li> </ol>

## StorNext Installation, Replication, HA, and Other Known Issues

The table below lists known issues specific to StorNext installations, data replication, HA systems, and other areas.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	68849	n/a	<p>After an appliance firmware upgrade, you might be unable to use previously functioning tape devices because the <code>lin_tape</code> device driver was automatically unloaded during the upgrade.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, rebuild the <code>lin_tape</code> device driver as shown in the following example:</p> <pre>rpm -e lin_taped rpm -e lin_tape rpmbuild --rebuild /root/lin_tape-1.76.06-1.src.rpm rpm -ivh /root/rpmbuild/RPMS/x86_64/lin_tape-1.76.06-1.x86_64.rpm rpm -ivh /root/lin_taped-1.76.0-rhel6.x86_64.rpm</pre>

Operating System	Change Request Number	Service Request Number	Description/Workaround
Linux	70282	n/a	<p>The <a href="#">StorNext Connect Utilization App Version 1</a> does not recognize the HGST ActiveScale™ P100 (Quantum Lattus P100) integrated object storage system, and does not incorporate capacity or data movement associated with the HGST ActiveScale™ P100 (Quantum Lattus P100) system as a target.</p> <p><b>i Note:</b> This does not otherwise affect the functionality of the HGST ActiveScale™ P100 (Quantum Lattus P100) integrated object storage system.</p> <p><b>Workaround:</b> There is currently no workaround for this issue. If you experience this issue, contact <a href="#">Quantum Technical Support</a>.</p>

# Contacting Quantum

## Contacts

For information about contacting Quantum, including Quantum office locations, go to:

<http://www.quantum.com/aboutus/contactus/index.aspx>

For further assistance, or for training opportunities, contact the Quantum Customer Support Center:

Region	Support Contact
North America	1-800-284-5101 (toll free) +1-720-249-5700
EMEA	+800-7826-8888 (toll free) +49 6131 324 185
Asia Pacific	+800-7826-8887 (toll free) +603-7953-3010

For worldwide support:

<http://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support>

## Comments

To provide comments or feedback about this document, or about other Quantum technical publications, send e-mail to:

[doc-comments@quantum.com](mailto:doc-comments@quantum.com)